

Distributed Newest Vertex Bisection

Martin Alkämper^a and Robert Klöforn^b

^aInstitut für Angewandte Analysis und Numerische Simulation,
 Fachbereich Mathematik, Universität Stuttgart, Pfaffenwaldring
 57, D-70569 Stuttgart, Germany,
www.ians.uni-stuttgart.de/nmh/,
alkaemper@ians.uni-stuttgart.de

^bInternational Research Institute of Stavanger, P. O. Box 8046,
 4068 Stavanger, Norway, <http://www.iris.no>,
robert.kloefkorn@iris.no

March 17, 2016

Abstract

Distributed adaptive conforming refinement requires multiple iterations of the serial refinement algorithm and global communication as the refinement can be propagated over several processor boundaries. We show bounds on the maximum number of iterations. The algorithm is implemented within the software package DUNE-ALUGRID.

Keywords: Adaptive method, mesh refinement, parallel, DUNE

1 Introduction

Conforming finite elements over conforming, unstructured, adaptive grids have been shown to behave very well for numerical simulations of diffusive processes ([7]). On the other hand new computer architectures demand parallelism in algorithms and grids to reach their full potential. For an adaptive, parallel, unstructured and conforming grid we need a parallel (or distributed) refinement strategy.

In this paper we analyze the distributed refinement strategy called Distributed Newest Vertex Bisection. It is the straightforward extension [11] of the serial Newest Vertex Bisection (NVB) introduced by Sewell [18] and we will show that the parallel overhead is bounded, in particular by a constant independent of the number of processors (Theorem 4.8). The Distributed NVB as described in this paper has been implemented in the open-source package DUNE-ALUGRID [1] which is a module for the DUNE software framework [3, 4]. Domain decomposition in combination with adaptivity requires load balancing to equidistribute the workload. In DUNE-ALUGRID this is done by equilibrating the number of cells belonging to each processor. We make the common assumption that

load balancing is done after the refinement algorithm is finished. Hence we will not consider its effect on the computational cost of the refinement algorithm. However, we will show that it is possible to implement the Distributed NVB using modern techniques of parallel computing such as communication hiding to achieve excellent strong scaling on a petascale super computer.

Most implementations of parallel adaptive grids use nonconforming hexahedral (or quadrilateral in 2 dimensions) cells with some mesh-balance to acquire a mesh grading required for stability estimates. Implementations of conforming adaptive parallel meshes are scarce, as especially in more than two dimensions the refinement propagation is non-trivial.

The Distributed NVB refinement strategy is also implemented in the toolbox AMDiS [21], and we will show, that the bound they give on the communication ($O(\log P)$, where P is the number of partitions in [21, Sec. 2.4]) is too weak for large P , unless the decomposition fulfills additional assumptions.

Another approach to parallel simplex refinement has been published by Rivara et al. [17]. In this work a parallel algorithm is introduced that produces an unstructured conforming mesh, which is not parallel in the domain decomposition sense, but instead the whole grid is known on each processor and the algorithm itself is executed in parallel. Furthermore, the refinement strategy differs, as it is not Newest Vertex Bisection(NVB), but Longest Edge Bisection which was introduced by Rivara in [16]. In [9] a fully distributed parallel Longest Edge Bisection is implemented based on the package DOLFIN and reasonable scaling results up to 1024 cores are presented, however, by sacrificing the theoretical backing of the adaptation algorithm.

The rest of the paper is structured as follows. First we introduce NVB with examples for 2-dimensional grids. Then the NVB refinement algorithm is extended to work in decomposed domains. Afterwards we analyze the Distributed NVB and show bounds on the parallel overhead which are reflected in the numerical experiments. The results hold for grids of any dimension unless stated otherwise.

2 Newest Vertex Bisection

In this section we shortly introduce NVB for conforming triangulations in two space dimensions. It was introduced by Sewell [18] and enhanced by Mitchell with a recursive refinement algorithm [13, 14]; compare also with [2, 12, 19, 20]. We follow the notation of [8].

2.1 Recurrent bisection of a simplex

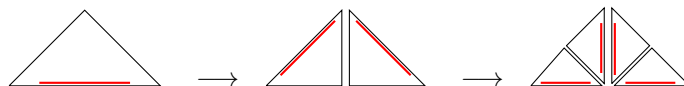


Figure 2.1: NVB: a triangle with its two children and four grandchildren. The refinement edges are indicated in red.

In order to easily describe NVB we identify a simplex T with its set of *ordered* vertices

$$T = [z_0, z_1, z_2].$$

The edge between the first and last vertex we call *refinement edge*. NVB refines T by inserting a new vertex in the midpoint $\bar{z} = \frac{1}{2}(z_0 + z_2)$ of the refinement edge $\overline{z_0 z_2}$ and

$$T_1 = [z_0, \bar{z}, z_1] \quad \text{and} \quad T_2 = [z_2, \bar{z}, z_1]$$

are the two children of T . This procedure automatically presets the children's refinement edges by the local ordering of their vertices. NVB thereby determines the refinement edge of any descendant produced by recurrent bisection of a given initial element T_0 from the vertex order of T_0 ; see Figure 2.1.

Recurrent bisection induces the structure of an *infinite binary tree* $\mathcal{F}(T_0)$: Any node T inside the tree is an element generated by recurrent application of NVB. The two successors of a node T are the children T_1, T_2 created by applying NVB to T .

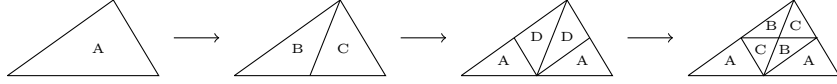


Figure 2.2: NVB: The four similarity classes for an initial element T_0 .

Finally, NVB produces shape regular descendants since all $T \in \mathcal{F}(T_0)$ belong to at most four similarity classes; compare with Figure 2.2. This is a consequence of the fact that NVB always bisects the angle at the newest vertex. In the end, any angle of any simplex is bisected at most once. This can be extended to any dimension d , see e.g. [19]. In the following, unless specified explicitly, we assume a general $d \geq 2$.

2.2 Recurrent refinement of triangulations with NVB

Let \mathcal{T}_0 be a conforming and exact triangulation of a bounded polygon $\Omega \subset \mathbb{R}^d$. We can refine \mathcal{T}_0 or a refinement \mathcal{T} of \mathcal{T}_0 by applying the NVB to selected simplices. More than the selected elements have to be refined when striking for conforming triangulations. Here we refer to [13] for a recursive refinement algorithm and to [2] for an iterative one.

We next introduce notations related to triangulations. The *master forest*

$$\mathcal{F} := \mathcal{F}(\mathcal{T}_0) = \bigcup_{T_0 \in \mathcal{T}_0} \mathcal{F}(T_0)$$

holds full information about all possible refinements of \mathcal{T}_0 . We denote by $\mathbb{T} = \mathbb{T}(\mathcal{T}_0)$ the class of all *conforming* refinements of \mathcal{T}_0 .

For $\mathcal{T} \in \mathbb{T}$ the sets of all its vertices and edges are \mathcal{V} and \mathcal{E} , respectively. For $T \in \mathcal{T}$ we set $\mathcal{V}(T) := \mathcal{V} \cap T$ and for $z \in \mathcal{V}$ we define $\mathcal{T}(z) := \{T \in \mathcal{T} \mid z \in T\}$. The finite element star at a vertex z is then $\Omega_z := \bigcup \{T : T \in \mathcal{T}(z)\}$. We let $h_{\mathcal{T}} \in L_{\infty}(\Omega)$ be the piecewise constant mesh-size function with $h_{\mathcal{T}|T} = h_T := |T|^{1/2} \approx \text{diam}(T)$ for $T \in \mathcal{T}$. We use $h_{\min, \max}(\mathcal{T})$ for the smallest and largest element size of \mathcal{T} . We say $T, T' \in \mathcal{T}$ are *direct neighbours* iff there is an $E \in \mathcal{E}$ with $E \subset T \cap T'$.

Important in the course of this article is the *generation* of an element. For each $T \in \mathcal{T}$ there is a $T_0 \in \mathcal{T}_0$ such that $T \in \mathcal{F}(T_0)$. The generation $\text{gen}(T)$ is the number of its ancestors in the tree $\mathcal{F}(T_0)$, or, equivalently, the number of bisections needed to create T from T_0 .

The following simple properties are useful.

Lemma 2.1. (1) *For $T \in \mathcal{F}(T_0)$ with $T_0 \in \mathcal{T}_0$ we have*

$$h_T = 2^{-\text{gen}(T)/2} h_{T_0}.$$

(2) *Defining $\alpha_0 := \max\{\#\mathcal{T}(z_0) \mid z_0 \in \mathcal{V}_0\}$ we have for $d = 2$ and $z \in \mathcal{V}$ the bound*

$$\#\mathcal{T}(z) \leq \begin{cases} 8 & \text{if } z \in \mathcal{V} \setminus \mathcal{V}_0, \\ 2\alpha_0 & \text{if } z \in \mathcal{V}_0. \end{cases}$$

Proof. Bisection halves the volume of a simplex. The definition of $\text{gen}(T)$ then gives the first claim. During refinement any angle is bisected at most once, which yields the second assertion. \square

The following assumption on a compatible distribution of refinement edges in \mathcal{T}_0 is instrumental in the analysis of NVB, like the complexity estimates in [5, 19, 10].

Assumption 2.2 (Compatibility Condition). Suppose $T, T' \in \mathcal{T}_0$ are direct neighbours with common edge $T \cap T' = E \in \mathcal{E}_0$. Then either E is the common refinement edge of both T and T' , or E is the refinement edge of descendants T'' of T and T''' of T' such that $\text{gen } T'' = \text{gen } T''' < d$.

Mitchell has shown that a distribution of refinement edges, s.t. assumption 2.2 holds, can be found for any initial triangulation \mathcal{T}_0 [13, Theorem 2.9] in 2 dimensions; compare also with [5, Lemma 2.1]. The assumption particularly implies that any uniform refinement of \mathcal{T}_0 is conforming, i.e., for any $g \in \mathbb{N}_0$ we find that $\{T \in \mathcal{F}(\mathcal{T}_0) \mid \text{gen}(T) = g\} \in \mathbb{T}$. The proof of this property is a combination of [20, §4] and [19, Theorem 4.3]. It is the key to show the following property of NVB; compare with [19, Corollary 4.6].

Proposition 2.3 (Characteristics of NVB). *Suppose that the initial triangulation \mathcal{T}_0 satisfies Assumption 2.2. Let $\mathcal{T} \in \mathbb{T}$ be given and suppose that $T, T' \in \mathcal{T}$ are direct neighbours such that the common edge $E = T \cap T'$ is the refinement edge of T . Then we either have $\text{gen}(T') = \text{gen}(T)$ and E is also the refinement edge of T' , or $\text{gen}(T') = \text{gen}(T) - i$ with $1 \leq i < d$.*

A simple consequence is $|\text{gen}(T) - \text{gen}(T')| \leq d - 1$ for direct neighbours $T, T' \in \mathcal{T}$.

3 Distributed Newest Vertex Bisection

The extension of NVB to the domain decomposition case is necessary, as the decomposed grid needs to be conforming across processor/partition boundaries. The basic idea is to execute the serial algorithm on each partition, communicate the refinement status of the partition boundary to the corresponding neighbour, refine conformingly and iterate.

In [11] it is shown that this parallel Algorithm 3.1 yields the same final triangulation as the serial version for meshes that fulfill the compatibility condition 2.2. This is essentially due to the fact that there is a unique mapping from the set of marked elements to the final refinement situation.

Algorithm 3.1: Distributed Newest Vertex Bisection

```

1 Initialize set of elements marked for refinement on each Partition  $M_i$ ,
   $0 \leq i < P$ .
2 while  $M_i \neq \emptyset \forall i$  do
3   Refine Partition  $P_i$  using NVB until  $M_i$  is empty
4   Communicate refinement status of partition boundary to
    corresponding neighbour
5   Add nonconforming simplices to  $M_i$ 
6   Communicate globally, whether  $M_i$  is empty.
7 end
```

We improved the algorithm introduced in [11] by additionally communicating the edge status (i.e. whether the edge has been bisected) of edges belonging to the process boundaries. This is slightly more communication expensive in 3 or more dimensions but it reduces the number of iterations needed. For 2 dimensions both algorithms coincide as faces are always 1 dimensional.

The stopping criteria of the while loop requires a global communication (Allreduce, $O(\log p)$ where p is the number of partitions), which cannot be expected to scale well onto many cores. On the other hand communicating the refinement status to the neighbour can be expected to scale quite well as long as the number of neighbouring partitions stays small, which relates to the quality of the decomposition.

4 Communication of the Distributed Newest Vertex Bisection

Bounds for the amount of communication necessary to reach a conforming triangulation are directly related to bounding the number of iterations in Algorithm 3.1. While the first bound from Theorem 4.6 does not need more assumptions than Compatibility Condition 2.2, Theorem 4.8 additionally requires dimension $d = 2$ and a certain form of mesh decomposition.

The following first lemma helps to understand the direct consequences of a single refinement.

Lemma 4.1. *For all direct neighbours T' of an element $T \in \mathcal{T}$ with refinement edge E with $E \subset T' \cap T$ one of the following two statements holds*

- 1 *Refinement of E in T' induces no further refinement (it already is the refinement edge)*
- 2 *E is the refinement edge of a child of T' and refinement of E induces up to $d - 1$ refinements of elements $T_i \subset T'$ with $\text{gen}(T_i) < \text{gen}(T)$.*

Proof. Two cases:

- 1 $\text{gen}(T) = \text{gen}(T') \Rightarrow$ no further refinement. It is the refinement edge of both elements because of the compatibility condition.

2 $\text{gen}(T) = \text{gen}(T') + i$ with $1 \leq i \leq d - 1 \Rightarrow$ induces i further refinements. The refinement edge can only be the same, if both elements have the same generation and direct neighbours can only differ in generation by d at most. The generation is increased by one with each refinement, so there have to be i refinements of descendants of T' that are refined, until the i -th descendant shares the refinement edge with T and yields the conforming closure of that element.

□

Lemma 4.1 can be applied recursively. A single refinement may lead to refinements of direct neighbours at 1 to $(d - 1)$ generations lower and these may again lead to refinements at even lower generations.

Example 4.2. Let us assume a simplex T with generation $\text{gen}(T) = 0$ and its direct neighbour T' with $\text{gen}(T') = d - 1$ and refinement edge $E = T \cap T'$. Now we refine T' , so we refine E and hence T . The compatibility condition 2.2 yields that E is the refinement edge of a descendant T'' of T with $\text{gen}(T'') = d - 1$. We denote by T_E^i the descendant of T with generation i that contains E and we denote its refinement edge by E_i . So $T'' = T_E^{d-1}$, $T = T_E^0$ and $E_{d-1} = E$. Then the Refinement Propagation can be depicted in the following graph of figure 4.1.

The dashed lines denote the direct closure. It cannot induce further refinement

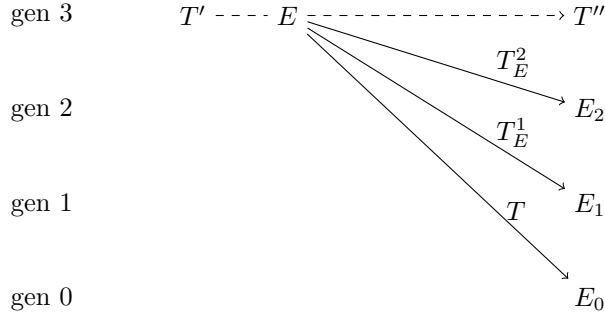


Figure 4.1: The direct Refinement Propagation of example 4.2 with $d = 4$.

and can be neglected. More general we have to analyze the direct Refinement Propagation for any element that contains E and additionally for all elements containing any of the edges E_i , as their refinement may also induce further refinement in the grid.

This leads us to the following definition.

Definition 4.3. Refinement of an element T with refinement edge e_0 and generation $\text{gen}(T) = l$ induces refinement propagation in form of a directed graph with root e_0 . For any element T' with $e_0 \subset T'$ and $\text{gen}(T') =: l' < l$ we have a directed edge from e_0 to the refinement edge E'_i of $T_{e_0}^{l'}$ for $l' \leq i < l$ as new nodes. We repeat by setting every newly introduced node as a local root. We call this the *Refinement Propagation Graph*.

Example 4.4. Figure 4.2 depicts an example of the Refinement Propagation graph of an initial refinement of an element T with refinement edge e_0 . For some elements the refinement results in the direct closure, so they are not included in

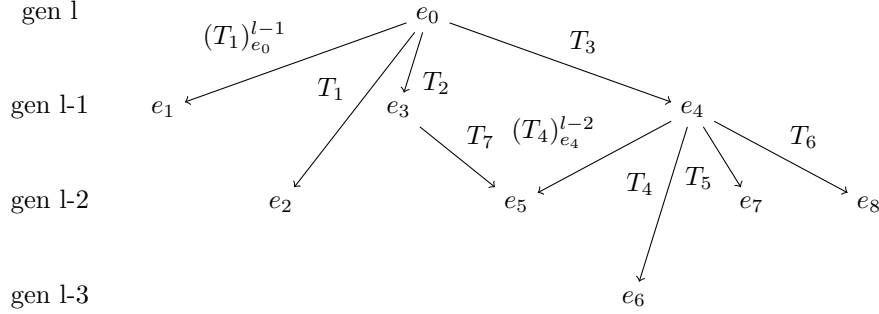


Figure 4.2: An example for an Refinement Propagation Graph.

the graph as well as T . For three direct neighbours T_1, T_2, T_3 of T the refinement does not result in the direct closure. For T_1 it even results in an additional refinement of its child $(T_1)_{e_0}^{l-1}$ that contains e_0 . Bisection of e_4 to refine T_3 is locally similar to the refinement of T by e_0 . Note that the graph is not a tree, as refinement edges may be shared (e_5 in our example). All leaves do not induce further refinement, so all adjacent elements are of the same level and refinement is the direct closure, which is not included in the graph.

Remark 4.5. In 2 dimensions the Refinement Propagation Path consists solely of nodes of degree 2 (and the root and the leaf), since in 2 dimensions every edge is shared by exactly two elements and the direct closure is not included. Hence in 2 dimensions we call it the *Refinement Propagation Path*.

With this preliminary work and exploiting the compatibility condition 2.2 we state the following theorem.

Theorem 4.6. *Let $M \subset \mathcal{T}$ be the set of elements marked for refinement. Then the number N of iterations in the Algorithm 3.1 to reach a conforming state satisfies*

$$N \leq \max_{T \in M} \max_{T' \in \mathcal{T}} (\text{gen}(T) - \text{gen}(T')) + 2.$$

Proof. Let $T \in M$ with $\text{gen}(T) = l$. We will bound the maximum depth of the Refinement Propagation Graph of T , which is an upper bound for the number of iterations as in the worst-case scenario refinement needs to be communicated at every edge.

Due to Lemma 4.1 refinement can be propagated at generation $l-d < \text{gen}(T') < l$, in particular at generation $l-1$ and no propagation at generation l . So the maximum number of propagations N_T resulting from refining T is $l - \min_{T' \in \mathcal{T}} \text{gen}(T')$. This is clearly also the maximum depth of the Refinement Propagation graph.

We have to take into account, that the Refinement Propagation graph does not consider the direct closure, which could need an additional communication. It follows

$$N_T \leq l - \min_{T' \in \mathcal{T}} \text{gen}(T') + 1$$

If we now take the maximum over all $T \in M$ this results in

$$\max_{T \in M} \max_{T' \in \mathcal{T}} (l(T) - l(T')) + 1.$$

We have to add another 1 as Algorithm 3.1 has to communicate that it has finished. \square

The following example demonstrates that this bound is sharp and that we cannot expect anything better even in the simple case of 2 partitions. In particular the bound $O(\log p)$ from [21, Sec. 2.4] cannot hold without further assumptions on the decomposition.

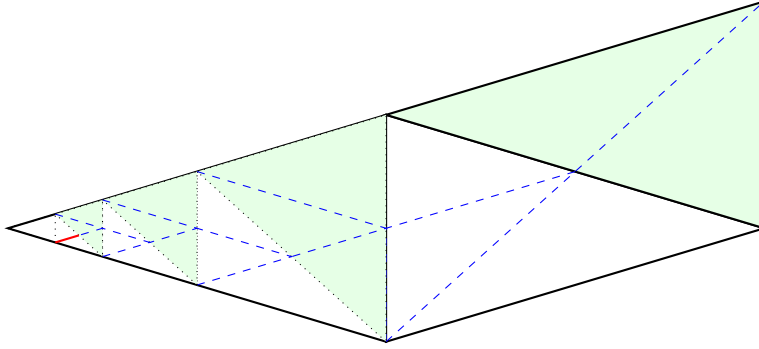


Figure 4.3: A distributed refined mesh to illustrate that the bound of theorem 4.6 is sharp. Partitions are indicated by the color of the cells. Black lines denote macro element borders. Dotted lines denote the initial refinement situation. The refinement request is marked in red. Blue dashed lines denote its Refinement Propagation.

Example 4.7. Figure 4.3 shows a mesh consisting of three initial triangles fulfilling the compatibility condition 2.2. One triangle has been refined into a corner, so the minimum generation in the mesh is $\min_{T' \in \mathcal{T}} \text{gen}(T') = 0$ and the generation of the marked element $T \in M = \{T\}$ is $\text{gen}(T) = 6$, so the bound from Theorem 4.6 is 8 iterations. Every refinement induces additional refinement at exactly one generation lower, so the refinement propagation path traverses 7 edges. The mesh is partitioned in such a way that every edge lies on a processor boundary, which implies that after every refinement Algorithm 3.1 has to stop and globally communicate. This means we get 7 iterations, where the marked set is not empty on all processors and one final iteration to communicate, that we are finished. In total this is 8 iterations, which is the bound predicted from Theorem 4.6.

Distributing the elements into partitions as depicted in figure 4.3 is not purely artificial. Sorting the leaf elements in vertical direction with respect to their center coordinates leads to these partitions.

Example 4.7 demonstrates that we have to impose additional assumptions on the decomposition to expect better bounds on the number of iterations of Algorithm 3.1.

A reasonable assumption could be that the partitions are created by a hierarchical space-filling curve, such that elements that are close in the refinement-tree are probably on the same partition.

Another assumption is that the mesh is partitioned solely on elements of a generation l^* which are then distributed onto partitions together with their

respective refinement trees. This second assumption will be analyzed in this paper, because this is the partitioning currently implemented in our grid manager DUNE-ALUGRID. Without loss of generality we can assume that we distribute elements on the macro level (i.e. elements T with $\text{gen}(T) = 0$) with their respective refinement trees. If this was not the case we could set the uniformly refined grid as our new initial grid, as for this kind of partitioning coarsening below generation l^* is forbidden.

The following results hold for partitioning on any fixed level but only for 2-dimensional grids, as we rely on the fact that we have a Refinement Propagation Path instead of a full graph.

Theorem 4.8. *Let dimension $d = 2$ and the mesh be partitioned as described above. Let $z \in \mathcal{V}$ and let $N_z = \{T \in \mathcal{T}_0 : z \subset T\}$ the set of macro elements containing that vertex. For an element T in the set of marked elements M let $T_0(T) \supset T$ be the element of the macro grid \mathcal{T}_0 that contains T . Then the number of global communications N in Algorithm 3.1 satisfies*

$$N \leq \max_{T \in M} \max_{z \in T_0(T)} \frac{3}{4} \#N_z + \frac{7}{4} \leq \max_{z \in \mathcal{T}_0} \frac{3}{4} \#N_z + \frac{7}{4}.$$

Proof. The second inequality is trivial. The proof of the first inequality splits into two parts.

1. Refinement propagation around vertex z .
2. Refinement propagation inside of macro elements that contain vertex z .

We start with a similar observation as in the proof of Theorem 4.6. By bounding the traversals of edges of \mathcal{T}_0 within the refinement propagation path, we bound the number of global communications. We count edges of the refinement propagation path that are contained in edges of \mathcal{T}_0 .

Let $T \in M$ with $z \in T_0(T)$ be the element to be refined. Refinement of T leads to elements of generation $l = \text{gen}(T) + 1$. All refinement propagation of refinement of T is contained in the refinement propagation of uniformly refining T_0 to this level, which we will investigate. (cf. Figure 4.4)

(1): Refinement propagation around vertex z .

We are now investigating the effect of refinement of T_0 at its vertex z . There are two possibilities:

- a. z is opposite of the initial refinement edge of T_0 .

Then there are two leaf elements $T^{0,1}$ with $z \in T^{0,1} \subset T_0$ and $\text{gen}(T^0) = \text{gen}(T^1)$. If $\text{gen}(T^{0,1})/2 \bmod 2 = 0$, the refinement edge of T^0 and T^1 is the shared edge. If $\text{gen}(T^{0,1})/2 \bmod 2 = 1$ the refinement edge of $T^{0,1}$ is a subedge of an edge of T_0 containing z and as we are uniformly refining up to level l , every odd level there is refinement propagation across these edges.

- b. z is contained in the initial refinement edge of T_0 .

Then there is one leaf element with $z \in T' \subset T_0$. The refinement edge of T' is always contained in one of the edges of T_0 . So every level refinement of T' leads to refinement propagation across one of the two edges of T_0 containing z .

We know that elements around the vertex z form the refinement propagation path of T' , as they all differ by one in generation and due to the argument above.

The path evolves around the vertex until it encounters an element with the lowest level. The maximum level difference around the vertex is $(\#\mathcal{T}(z) - 1)/2$ as there have to exist at least two elements with lowest generation. So the maximum length of the refinement propagation path around z is $(\#\mathcal{T}(z) - 1)/2$. We want to bound the refinement propagation path edge traversals with respect to $\#N_z$. $\#N_z$ is smaller than $\#\mathcal{T}(z)$, as for every element in $\#N_z$ where the refinement edge is opposite of z , there are two elements in $\#\mathcal{T}(z)$. Now there are two cases:

a. $\#N_z$ is even:

The worst case is $\#\mathcal{T}(z) = 3/2\#N_z$ and all refined macro angles are neighbouring. Then, in one of the circumvention direction all edge traversals are traversals in N_z and so we may need $(3/2\#N_z - 1)/2$ traversals in N_z until we encounter the element with the lowest level.

b. $\#N_z$ is odd:

The worst case is $\#\mathcal{T}(z) = 3/2\#N_z + 1/2$ and all refined angles are neighbouring. Then, in one of the circumvention direction all edge traversals are traversals in N_z and so we may need $(3/2(\#N_z) + 1/2 - 1)/2$ traversals in N_z until we encounter the element with the lowest level.

So the number of macro edge traversals N_z^e within the refinement propagation path around z satisfies

$$N_z^e \leq (3/2(\#N_z - 1/2))/2 = 3/4\#N_z - 1/4.$$

(2): Refinement propagation inside of macro elements that contain vertex z .

We already know that the uniform refinement propagates into all macro elements, which share z . So we can neglect macro elements that share an edge with T_0 as we know that their refinement propagation does not yield any additional information.

So we consider a macro element T_1 that contains z and does not share an edge with T_0 . We know that we have a refinement at vertex z and want to investigate, whether we can reach the edge opposite of z . The other edge does not matter, as it contains z .

The size of an element of generation l is $2^{-l}|T_0|$ (Lemma 2.1). The refinement propagation path consists of elements that differ by one in level. So the size of the refinement propagation path inside the element T_1 is $\sum_k 2^{-k}|T_1| = |T_1|(2^{-n} - 2^{-l})$. So to get to another element from the opposite vertex it has to include an element of level 0, a macro element. This is only possible, if the initial refinement edge is opposite of z .

In combination with the previous result this yields

$$N \leq (3/2\#N_z - 1/2)/2 + 1 = 3/4\#N_z + 3/4.$$

Now we finish the proof by adding another $+1$ for communicating the final status and taking the maximum over all vertices of T_0 and over all marked elements. \square

Remark 4.9. From part (2) of the proof one can see, that if the mesh is uniformly refined up to one level below the macro level, then the $+1$ from this part disappears and we get.

$$N \leq \max_{T \in \mathcal{M}} \max_{z \in T_0(T)} 3/4\#N_z + 3/4 \leq \max_{z \in T_0} 3/4\#N_z + 3/4$$

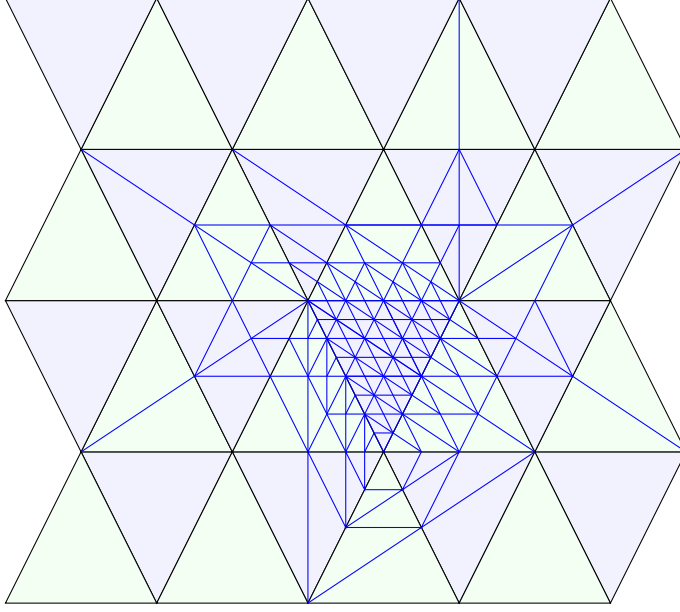


Figure 4.4: Refinement propagation from uniform refinement of a single macro element.

Remark 4.10. If we do not count all edges of \mathcal{T}_0 , but only those that are actually processor boundaries, we get the following bound that is better as long as every processor gets "nice" partitions.

$$N \leq \max_{T \in M} \max_{z \in T_0(T)} \#P_z - 1 + 1 = \max_{T \in M} \max_{z \in T_0(T)} \#P_z$$

where $\#P_z$ is now the number of partitions that share a vertex. The proof is similar to the proof of theorem 4.8, but we get $\#P_z - 1$ as we cannot argue with circumventions in both directions and the $+1$ is again due to the final communication. Note that if a processors partition p has several elements containing z and there is no path of elements $T \in p$ containing z connecting two elements, every connected subdomain has to be counted as a partition.

Remark 4.11. Theorem 4.8 provides a mesh constant. So no dependence on number of processors is required, but just a regular initial mesh (see α_0 in Lemma 2.1).

Remark 4.12. We believe that theorem 4.8 holds in a similar way for higher dimensions. It is a hard problem as the shape of the Refinement Propagation graph is not known.

Based on the estimate in remark 4.10 we propose a new improved algorithm for

compatible meshes.

Algorithm 4.1: Improved Distributed Newest Vertex Bisection

```

1 Initialize set of elements marked for refinement on each Partition  $M_i$ ,
   $0 \leq i < P$ .
2 for  $i = 0, \dots, \max_{z \in \mathcal{T}_0} \#P_z - 1$  do
3   Refine Partition  $P_i$  using NVB until  $M_i$  is empty
4   Communicate refinement status of partition boundary to
    corresponding neighbour
5   Add nonconforming simplices to  $M_i$ 
6 end
```

For compatible meshes we have proven, that this algorithm yields the conforming closure and reaches the final status. So communicating the final status is no longer necessary. Hence we take the better bound $\max_{z \in \mathcal{T}_0} \#P_z - 1$ instead of $\max_{z \in \mathcal{T}_0} \#P_z$. We can directly use the bound from Theorem 4.8 and get an algorithm, which does not need global communication at all. Unfortunately we cannot expect meshes to be compatible, especially in 3 dimensions. In this case we propose to use a mixture of both algorithms, where a fixed number of loops is done like in 4.1 before switching to the Algorithm 3.1 after a global communication, whether $M_i \neq \emptyset \forall i$.

5 Implementation

The NVB algorithm is implemented in the open-source package DUNE-ALUGRID available at <https://gitlab.dune-project.org/extensions/dune-alugrid>. [1] contains a description of the software and various examples.

In this section we discuss the implementation of the communication procedures which is not contained in detail in [1]. Communication is needed to exchange the refinement flags, e.g. line 4 in Algorithm 3.1 and 4.1. It is essential for the Distributed NVB that this is done in a very efficient way to guarantee excellent scalability. In DUNE-ALUGRID we have chosen to interleave the send and receive procedures with the packing and unpacking of refinement information. This way we are able to hide some of the communication latency behind the necessary pack and unpack of information. We briefly sketch the send and pack routine as well as the receive and unpack routine used in DUNE-ALUGRID.

Let \mathcal{L}_p^s be the set of all ranks that process $p \in [0, P - 1]$ sends data to and \mathcal{L}_p^r the corresponding set p received messages from. We call the communication symmetric if $\mathcal{L}_p^s = \mathcal{L}_p^r$. Asymmetric communication occurs, for example, during the load balancing, where the list of send and receive ranks can differ. The communication algorithm, however, is the same. Using the sets $\mathcal{L}_p^s, \mathcal{L}_p^r$ the corresponding methods for *pack-and-send* and *receive-and-unpack* are briefly explained in Algorithm 5.1 and 5.2, respectively. In the following \mathcal{T}_p^q denotes the set of simplices on process p with linkage to rank $q \in \mathcal{L}_p^{s,r}$.

Both algorithms are implemented in DUNE-ALUGRID and work for very general data sets and are thus used for all point to point communications in the package.

Algorithm 5.1: Pack and send

```
1 for  $q \in \mathcal{L}_p^s$  do
2   for  $T \in \mathcal{T}_p^q$  do
3     pack refinement information for simplex  $T$  and communication
      link  $q$ 
4   end
5   post non-blocking MPI_Isend for communication link  $q$ 
6 end
```

Algorithm 5.2: Receive and unpack

```
1 for  $q \in \mathcal{L}_p^r$  do
2    $r_q \leftarrow 0$ 
3 end
4  $n_r \leftarrow 0$ 
5 while  $n_r < |\mathcal{L}_p^r|$  do
6   for  $q \in \mathcal{L}_p^r$  do
7     if  $r_q = 0$  then
8        $r_q \leftarrow \text{MPI_Iprobe}(q)$ 
9       if  $r_q = 1$  then
10         $s \leftarrow \text{MPI_Getcount}(q)$ 
11        resize buffer for received message size  $s$ 
12        post MPI_Recv(  $q$  ) to write message from  $q$  to buffer
13        for  $T \in \mathcal{T}_p^q$  do
14          unpackData(  $T, q$  )
15        end
16         $n_r \leftarrow n_r + 1$ 
17      end
18    end
19  end
20 end
21 MPI_Waitall(  $\mathcal{L}_p^s$  ), see Algorithm 5.1
```

6 Numerical Experiments

In this section we show for various examples that the theoretical results can be reproduced and that very good scalability for the adaptation algorithm is observed on a petascale super computer.

6.1 Verification of theoretical results

The first experiment aims to reflect the theoretical results as we construct a worst-case experiment. A mesh is partitioned such that each process gets exactly one macro element. Then we refine a single element at one of its vertices up to level 20 and we examine the number of iterations necessary to reach the conforming status.

From figure 6.1 we see that for compatible 2d grids both bounds (red lines) are

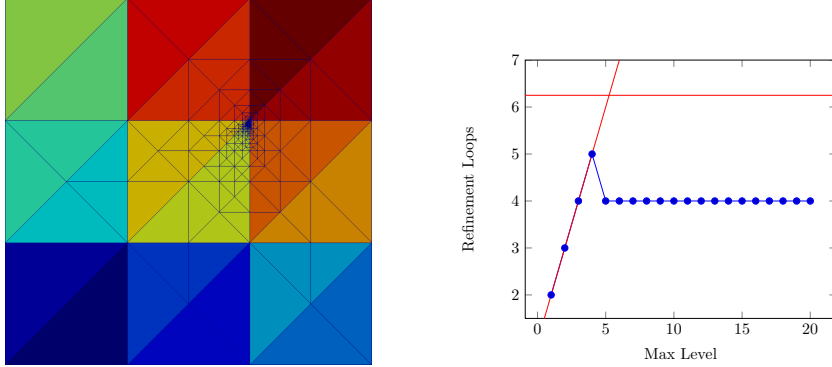


Figure 6.1: A 2d unit square with 18 macro elements on 18 processors. The central yellow element gets refined at the top vertex. The number of macro neighbours $\#N_z = 6$. In the right plot red lines denote the two bounds from the Theorems 4.8 and 4.6.

not violated by the current implementation. As this is a worst-case scenario, in the average case the behaviour is better. Now we perform the same experiment on a non-compatible grid.

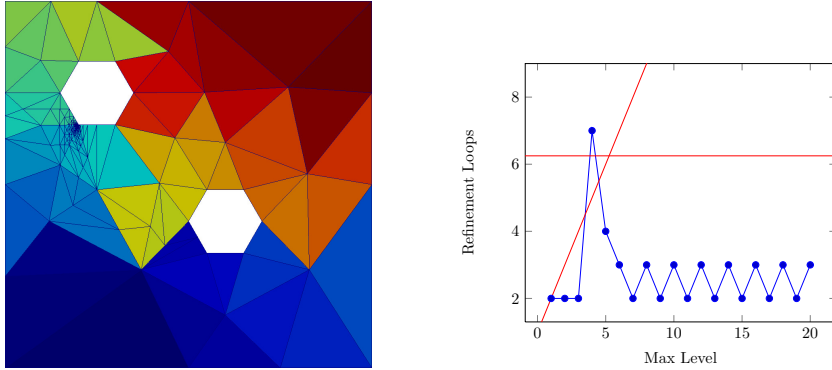


Figure 6.2: A non-compatible 2d grid with 60 macro elements on 60 processors. An element (with $\#N_z = 6$) gets refined at a vertex. In the right plot red lines denote the two bounds from the Theorems 4.8 and 4.6.

As expected, the plot in Figure 6.2 illustrates that compatibility is a necessary condition for both theorems. This means although the implementation is capable of handling non-compatible 2d grids, the bounds from the theory do not hold.

We expect the bound from Theorem 4.6 to hold in any dimension. This is not the case in figure 6.3. This failure arises from an implementation detail, that refinement status is communicated for faces instead of edges. After implementation of an additional communication of edges statuses during refinement we see that the theoretical result holds. In addition, the plot indicates the existence of a constant to bound the number of iterations in 3d similar to Theorem 4.8.

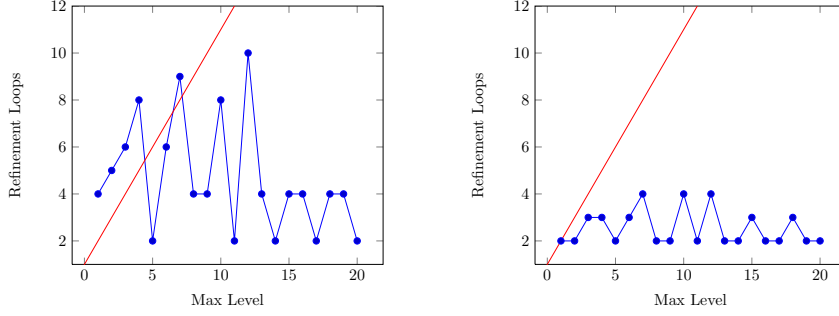


Figure 6.3: Unit cube. 162 Macro Elements/ 162 processors ($3 \times 3 \times 3$ Kuhn-dice). Central element gets refined at one vertex. On the left refinement status is communicated on faces. On the right refinement status is first communicated over edges and then additionally over faces.

6.2 Strong scaling experiments

In Figure 6.4 we show the refinement of a doughnut that rotates around the center (doughnut refinement). Triangles inside the doughnut are refined and triangles outside are coarsened. This test was introduced in [1] and serves as an excellent test for the Distributed NVB since frequent refinement and coarsening occurs throughout the simulation. In fact, the adaptation and load balancing is performed in each time step. The domain decomposition is based on the Hilbert space filling curve approach implemented in Zoltan [6]

Figures 6.5 and 6.6 we provide strong scaling results obtained for the doughnut refinement test in 2d and 3d, respectively. The scaling experiment have been performed on the super computer Yellowstone [15]. The observed strong scaling for the adaptation cycle (green triangular line) is excellent for all experiments and very close to the optimal scaling. The load balancing (pink boxed line) on the other hand at some points fails to scale well because the number of macro simplices per core becomes too small as the number of processes grows which results from the drawback of basing the partitioning upon the macro mesh. This is currently under investigation and will be reported in a separate article. In contrast to [21, Fig. 8], where for a different adaptation experiment the *mesh adaptation loop* yielded non-optimal strong scaling, we conclude from our investigations that the Distributed NVB scales well and depending on the macro mesh a fixed or very limited number global communications is needed.

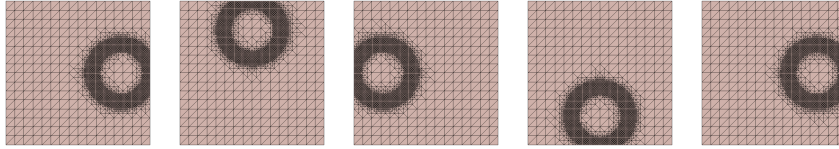


Figure 6.4: Refinement and coarsening of a doughnut like area that rotates around the center. From left to right the simulation time is increased from $t = 0$ to $t = 1$ by $\Delta t = 0.25$.

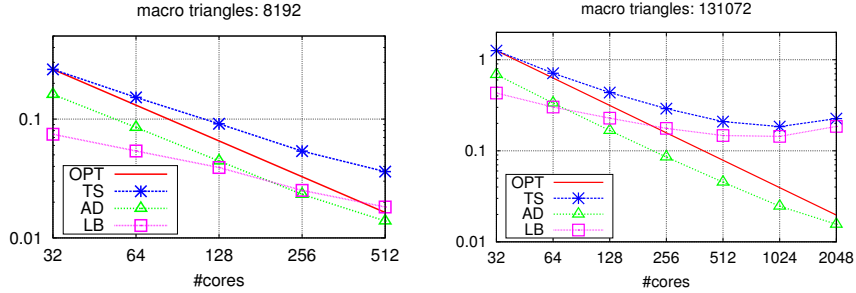


Figure 6.5: 2d results for the doughnut refinement test on a coarse triangular macro mesh (left) and a finer triangular macro mesh (right). For different number of cores the graph shows average run time per timestep in seconds for the different parts of the algorithm: a full time step (TS), the adaptation loop (AD), the load balancing (LB), and the expected optimal scaling (OPT). The scaling study has been performed on Yellowstone [15]. The test case is part of the DUNE-ALUGRID code and described in [1].

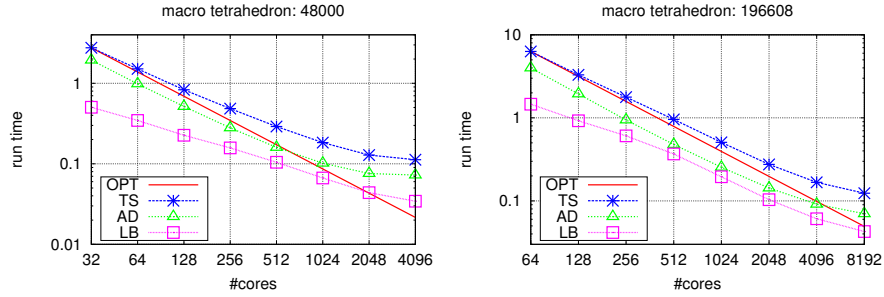


Figure 6.6: 3d results for the doughnut refinement test on a coarse tetrahedral macro mesh (left) and a finer tetrahedral macro mesh (right). For different number of cores the graph shows average run time per timestep in seconds for the different parts of the algorithm: a full time step (TS), the adaptation loop (AD), the load balancing (LB), and the expected optimal scaling (OPT). The scaling study has been performed on Yellowstone [15]. The test case is part of the DUNE-ALUGRID code and described in [1].

7 Summary

We have shown (and proven in 2d), that the number of iterations in Algorithm 3.1 to reach a conforming situation is bounded. In particular for grid implementations that do not partition the mesh on the leaf level, but on a certain fixed level, the bound is constant and independent of the current refinement situation. As the purpose of conforming grids is usually solving elliptic equations, the total run time is usually dominated by solving the equation. The performed *worst-case* experiments are reflected by the theory. These experiments also prove that the compatibility condition in Assumption 2.2 is essential and cannot be neglected. In addition, we presented a state of the art implementation of the Distributed NVB including asynchronous communication which is

needed to achieve excellent scaling on a petascale super computer. As a next step we will improve the flexibility of the load balancing algorithm which currently only allows to partition the macro mesh. For certain problems where singularities might arise, this might not be sufficient and yield poor scalability.

Acknowledgements

Martin Alkämper acknowledges the Cluster of Excellence in Simulation Technology (SimTech) at the University of Stuttgart for financial support. Robert Klöfkorn acknowledges NCAR/CISL’s Research and Supercomputing Visitor Program (RSVP) and the Research Council of Norway and the industry partners – ConocoPhillips Skandinavia AS, BP Norge AS, Det Norske Oljeselskap AS, Eni Norge AS, Maersk Oil Norway AS, DONG Energy A/S, Denmark, Statoil Petroleum AS, ENGIE E&P NORGE AS, Lundin Norway AS, Halliburton AS, Schlumberger Norge AS, Wintershall Norge AS – of The National IOR Centre of Norway for financial support.

References

- [1] ALKÄMPER, M., DEDNER, A., KLÖFKORN, R., AND NOLTE, M. The DUNE-ALUGrid Module. *Archive of Numerical Software* 4, 1 (2016), 1–28.
- [2] BÄNSCH, E. Local mesh refinement in 2 and 3 dimensions. *IMPACT of Computing in Science and Engineering* 3, 3 (1991), 181 – 191.
- [3] BASTIAN, P., BLATT, M., DEDNER, A., ENGWER, C., KLÖFKORN, R., KORNUBER, R., OHLBERGER, M., AND SANDER, O. A Generic Grid Interface for Parallel and Adaptive Scientific Computing. Part II: Implementation and Tests in DUNE. *Computing* 82, 2–3 (2008), 121–138.
- [4] BASTIAN, P., BLATT, M., DEDNER, A., ENGWER, C., KLÖFKORN, R., OHLBERGER, M., AND SANDER, O. A Generic Grid Interface for Parallel and Adaptive Scientific Computing. Part I: Abstract Framework. *Computing* 82, 2–3 (2008), 103–119.
- [5] BINEV, P., DAHMEN, W., AND DEVORE, R. Adaptive Finite Element Methods with convergence rates. *Numerische Mathematik* 97, 2 (2004), 219–268.
- [6] BOMAN, E. G., CATALYUREK, U. V., CHEVALIER, C., AND DEVINE, K. D. The Zoltan and Isorropia parallel toolkits for combinatorial scientific computing: Partitioning, ordering, and coloring. *Scientific Programming* 20, 2 (2012).
- [7] BONITO, A., AND NOCHETTO, R. H. Quasi-optimal convergence rate of an adaptive discontinuous Galerkin method. *SIAM J. Numer. Anal.* 48, 2 (2010), 734–771.

- [8] GASPOZ, F. D., HEINE, C.-J., AND SIEBERT, K. G. Optimal grading of the newest vertex bisection and H^1 -stability of the L^2 -projection. *IMA Journal of Numerical Analysis* (2015).
- [9] JANSSON, N., HOFFMAN, J., AND JANSSON, J. Framework for massively parallel adaptive finite element computational fluid dynamics on tetrahedral meshes. *SIAM J. Sci. Comput.* *34*, 1 (Feb. 2012), 24–41.
- [10] KARKULIK, M., PAVLICEK, D., AND PRAETORIUS, D. On 2D Newest Vertex Bisection: Optimality of Mesh-Closure and H^1 -Stability of L^2 -Projection. *Constructive Approximation* *38*, 2 (2013), 213–234.
- [11] LIU, Q., MO, Z., AND ZHANG, L. A parallel adaptive finite-element package based on ALBERTA. *International Journal of Computer Mathematics* *85*, 12 (2008), 1793–1805.
- [12] MAUBACH, J. Local Bisection Refinement for N-Simplicial Grids Generated by Reflection. *SIAM Journal on Scientific Computing* *16*, 1 (1995), 210–227.
- [13] MITCHELL, W. F. Unified multilevel adaptive finite element methods for elliptic problems. Phd thesis, University of Illinois, Urbana, IL,, 1988.
- [14] MITCHELL, W. F. A comparison of adaptive refinement techniques for elliptic problems. *ACM Trans. Math. Softw.* *15*, 4 (1989), 326–347.
- [15] NCAR/CISL. Computational and Information Systems Laboratory. Yellowstone: IBM iDataPlex System (Climate Simulation Laboratory). Boulder, CO: National Center for Atmospheric Research, 2012.
- [16] RIVARA, M.-C. Mesh Refinement Processes Based on the Generalized Bisection of Simplices. *SIAM Journal on Numerical Analysis* *21*, 3 (1984), 604–613.
- [17] RIVARA, M.-C., CALDERON, C., FEDOROV, A., AND CHRISOCHOIDES, N. Parallel decoupled terminal-edge bisection method for 3D mesh generation. *Engineering with Computers* *22*, 2 (2006), 111–119.
- [18] SEWELL, E. Automatic Generation of Triangulations for Piecewise Polynomial Approximation. Phd thesis, Purdue University, 1972.
- [19] STEVENSON, R. The completion of locally refined simplicial partitions created by bisection. *Math. Comput.* *77*, 261 (2008), 227–241.
- [20] TRAXLER, C. An algorithm for adaptive mesh refinement in n dimensions. *Computing* *59*, 2 (1997), 115–137.
- [21] WITKOWSKI, T., LING, S., PRAETORIUS, S., AND VOIGT, A. Software concepts and numerical algorithms for a scalable adaptive parallel finite element method. *Advances in Computational Mathematics* (2015), 1–33.